arXiv:1610.10061v1 [cs.DC] 31 Oct 2016

A GPU-Based Genetic Algorithm for the P-Median Problem

Bader F. AlBdaiwi and Hosam M.F. AboElFotoh

Computer Science Department, Kuwait University, Kuwait

{bdaiwi, hosam}@cs.ku.edu.kw

November 1, 2016

Abstract

The p-median problem is a well-known NP-hard problem. Many heuristics have been proposed in the literature for this problem. In this paper, we exploit a GPGPU parallel computing platform to present a new genetic algorithm implemented in Cuda and based on a Pseudo Boolean formulation of the p-median problem. We have tested the effectiveness of our algorithm using a Tesla K40 (2880 Cuda cores) on 290 different benchmark instances obtained from OR-Library, discrete location problems benchmark library, and benchmarks introduced in recent publications. The algorithm succeeded in finding optimal solutions for all instances except for two OR-library instances, namely pmed30 and pmed40, where better than 99.9% approximations were obtained.

Keywords P-Median Problem; NP-Hard; GPGPU; Cuda; Pseudo Boolean Formulation; Genetic Algorithms; Heuristics.

1 Introduction

The P-Median Problem (PMP) is formally defined as follows. Given a set $C = \{1, \ldots, n\}$ of n clients, a set $F = \{1, \ldots, m\}$ of m facilities, an integer $p < m$, and the distance d_{ij} between client i, $1 \leq i \leq n$, and facility j, $1 \leq j \leq m$. Let $y_{ij} \in \{0, 1\}$ be a decision variable such that $y_{ij} = 1$ if and only if client i is serviced by facility j, and let $x_j \in \{0, 1\}$ be a decision variable such that $x_j = 1$ if and only if facility j is open for service. The PMP objective is to minimize the total distance[1]

$$f_C(x, y) = \sum_{i \epsilon C} \sum_{j \epsilon F} x_j y_{ij} d_{ij}, \tag{1}$$

[1] We shall use distance and cost interchangeably

subject to

$$\sum_{j \in F} y_{ij} = 1 \quad \forall i \in C, \tag{2}$$

$$\sum_{j \epsilon F} x_j = p. \tag{3}$$

The objective function (1) minimizes the total distance between clients and the corresponding service facilities. Constraint (2) states that each client is serviced by exactly one facility. Constraint (3) states that the number of open facilities is exactly p. Let the set of open facilities be $O = \{o_1, \ldots, o_p\}$. Naturally, if client i is serviced by facility j, ($y_{ij} = 1$), then: 1) $j \in O$ (is open), and 2) d_{ij} is a minimum over $\{ d_{io_1} ... d_{io_p}\}$. An instance of the PMP is described by an $n \times m$ distance matrix $C = [d_{ij}]$ and a positive integer $p < m$. Note that we assume the elements of C are non-negative.

The PMP has a wide range of applications. It has been extensively researched in the literature. It has many applications in logistics [7][20] and location science [25][34]. It also has applications in finance and market analysis [15]. Unfortunately, it is NP-hard and hence difficult to solve it for optimality [23]. Comprehensive surveys on solving methods for the PMP and its variations can be referred to in [11][14][30][33].

In this paper, we exploit a GPGPU parallel computing platform to present a new genetic algorithm implemented in Cuda C version 7.5 (Compute Unified Device Architecture) and based on a Pseudo Boolean formulation of the PMP.

The rest of the paper is organized as follows. Section 2 introduces preliminaries and related literature review on the pseudo boolean formulation of the PMP, GPGPU and Cuda, and genetic algorithms. Section 3 presents the new algorithm. Section 4 highlights some implementation details. The algorithm time complexity is analyzed in Section 5. Section 6 presents the experimentation results, and Section 7 concludes the paper.

2 Preliminaries

2.1 A Pseudo Boolean Formulation of the PMP

The pseudo Boolean formulation of the PMP appeared in [18]. It is obtained as follows. For each client i, let $\prod^i = (\pi_{i1}, \ldots, \pi_{im})$ be an ordering of $1, \ldots, m$ such that $d_{i\pi_{ik}} \leq d_{i\pi_{il}}$ if $k < l$ for all $k, l \in \{1, \ldots, m\}$, and let $\triangle^i = (\delta_{i1}, \ldots, \delta_{im})$, where $\delta_{i1} = d_{i\pi_{i1}}$, and $\delta_{ir} = d_{i\pi_{ir}} - d_{i\pi_{i(r-1)}}$, for $r = 2, ..., m$. Therefore, the distance between client i, $i \in \{1, \ldots n\}$, and the facility serving it can be expressed using the following pseudo Boolean polynomial

$$d_i = \delta_{i1} + \sum_{k=2}^{m} \delta_{ik} \prod_{r=1}^{k-1} \bar{x}_{\pi_{ir}} \ . \tag{4}$$

Thus, the PMP can be reformulated as: Given $\triangle = [\delta_{ij}]$, $\prod = [\pi_{ij}]$, and p, find an assignment in $\{0, 1\}$ to x_i, $i \in \{1, \ldots, n\}$, such that

$$\sum_{j \in F} x_j = p, \tag{5}$$

and

$$B_C(z) = \sum_{i=1}^{n} \left(\delta_{i1} + \sum_{k=2}^{m} \delta_{ik} \prod_{r=1}^{k-1} z_{\pi_{ir}} \right) \tag{6}$$

is minimized. The Boolean variable $z_j = \bar{x_j}$ is 1 iff $x_j = 0$, denoting a closed facility. Note that the k^{th} term in Equation (4) contains $\prod_{r=1}^{k-1} \bar{x}_{\pi_{ir}}$. Therefore, this term must be zero $\forall\ k > m-p+1$ since at least one facility (say f) is open in any $m-p+1$ locations resulting in $\bar{x}_f = 0$. Thus, $\prod$ and $\triangle$ can be reduced to $\prod^{'}$ and $\triangle^{'}$ by omitting the last $p-1$ columns.

The objective function based on $\prod^{'}$ and $\triangle^{'}$ is known as *Hammer-Bersnev polynomial* (*HBP*)[5]. It can be further reduced through monomial reduction. Interested readers could refer to [4] for details. Example 1 illustrates the PMP pseudo boolean formulation.

Example 1 : Consider a PMP instance with $n = 5$, $m = 4$, $p = 2$ and

$$C = \begin{bmatrix} 7 & 10 & 16 & 11 \\ 15 & 17 & 7 & 7 \\ 10 & 4 & 6 & 6 \\ 7 & 11 & 18 & 12 \\ 10 & 22 & 14 & 8 \end{bmatrix}.$$

An ordering matrix $\prod$ and the corresponding matrix $\triangle$ are given by

$$\prod = \begin{bmatrix} 1 & 2 & 4 & 3 \\ 3 & 4 & 1 & 2 \\ 2 & 3 & 4 & 1 \\ 1 & 2 & 4 & 3 \\ 4 & 1 & 3 & 2 \end{bmatrix}, \ \triangle = \begin{bmatrix} 7 & 3 & 1 & 5 \\ 7 & 0 & 8 & 2 \\ 4 & 2 & 0 & 4 \\ 7 & 4 & 1 & 6 \\ 8 & 2 & 4 & 8 \end{bmatrix}.$$

Omitting the last $(p - 1 = 1)$ column corresponding to zero terms in Equation (4) results in:

$$\prod^{'} = \begin{bmatrix} 1 & 2 & 4 \\ 3 & 4 & 1 \\ 2 & 3 & 4 \\ 1 & 2 & 4 \\ 4 & 1 & 3 \end{bmatrix}, \triangle^{'} = \begin{bmatrix} 7 & 3 & 1 \\ 7 & 0 & 8 \\ 4 & 2 & 0 \\ 7 & 4 & 1 \\ 8 & 2 & 4 \end{bmatrix}.$$

The corresponding HBP representing total distance (cost) is

$$B_C(z) = \begin{array}{l}[7+3z_1+1z_1z_2]+[7+0z_3+8z_3z_4]+[4+2z_2+0z_2z_3]+\\ [7+4z_1+1z_1z_2]+[8+2z_4+4z_1z_4]\,.\end{array}$$

$B_C(z)$ has $n \times (m-p+1) = 15$ entries, and the reduced polynomial is

$$B_C(z) \;=\; 33+7z_1+2z_2+2z_4+2z_1z_2+8z_3z_4+4z_1z_4.$$

2.2 GPGPU and Cuda

Similar to many NP-hard problems, many heuristics have been developed for the PMP. Some of these heuristics tried to exploit parallel computing platforms to reach a near-optimal solution in a reasonably short time [11]. A few years ago, Nvidia has introduced Cuda (Compute Unified Device Architecture) that provides an application interface (API) for general purpose computing [**?**]. Hence, GPGPU refers to General Purpose computing on Graphics Processing Units. Nvidia graphics cards (as well as all graphics cards) are designed to do similar computations on large numbers of pixels. Therefore, they contain hundreds of processing elements (cores), although not as powerful as CPU cores, that execute thousands of similar threads (grouped in blocks) in parallel. Unlike CPU's threads, the context switching between blocks of threads requires a minimal overhead. There are not many GPU-based solutions for the PMP so far. Lim and Ma introduced GPU implementations for solving the PMP using the vertex substitution and Llod algorithms in [26] and [27]. Cuda C can be used for developing applications on GPGPU. In our implementation, we used Cuda C version 7.5.

A Cuda C program consists of two types of code: *Host* code and *Device* code. The Host code refers to that executed by the CPU and the device code refers to that executed by the GPU card. The Host code launches a *kernel* that is executed by the device. A kernel launch specifies the number of threads to be executed in parallel. These threads are grouped in blocks. Blocks in their turns are organized in a grid. All threads execute the same code but on multiple data, which Nvidia calls a *Single Instruction Multiple Thread* (SIMT) architecture. To support multidimensional data modeling and processing, Cuda enables defining grids and blocks to be single, double, or triple dimensions. For example, the following Cuda code declares two triple dimensional arrays using the *dim3* data type. Then, it invokes *KernelX* with a $3 \times 4 \times 6$ grid each element of which is a $2 \times 4 \times 4$ block.

$dim3\ Grid(3,4,6);$
$dim3\ Block(2,4,4);$
$KernelX <<< Grid, Block >>> (/*\ Parameter\ List\ */);$

The total number of blocks in *KernelX* is $(3 \times 4 \times 6 = 72)$, and each block has $(2 \times 4 \times 4 = 32)$ threads. Thus, the total number of threads in *KernelX* is $(72 \times 32 = 2,304)$. To differentiate among threads, Cuda defines two built-in

variables for block and thread indexing, namely *blockIdx* and *threadIdx*. Each of these variables is *3-dimensional*. For instance, *threadIdx* has three components *threadIdx.x, threadIdx.y*, and *threadIdx.z*. Cuda also defines two other built-in *3-dimensional* variables for the grid and block dimensions, *gridDim* and *blockDim*. They are automatically initialized at a kernel launch. In the above example, *KernelX* sets $gridDim.x = 3$, $gridDim.y = 4$, and $gridDim.z = 6$. It also sets $blockDim.x = 2$, $blockDim.y = 4$, and $blockDim.z = 4$.

Linearized unique block identifier, *BID*, and linearized unique global thread identifier, *TID*, can be derived from *gridDim*, *blockDim*, *blockIdx*, and *threadIdx* as follows:

$$\begin{aligned} BID &= blockIdx.x + blockIdx.y * gridDim.x + \\ &\quad gridDim.x * gridDim.y * blockIdx.z; \end{aligned}$$

$$\begin{aligned} TID &= BID * (blockDim.x * blockDim.y * blockDim.z) + \\ &\quad (threadIdx.z * (blockDim.x * blockDim.y)) + \\ &\quad (threadIdx.y * blockDim.x) + threadIdx.x; \end{aligned}$$

Threads may access data in parallel from different memory locations. Usually, the data accessed by each thread is determined by its index (within the block) or its global identifier. In Cuda, the device code can access only the device memory. Therefore, the host code has to initialize the device memory through Cuda calls that allocate device memory (*cudaMalloc*) and copy (*cudaMemcpy*) data between host RAM and device RAM. The device memory has three basic types:

Global can be accessed by all threads from all blocks.

Shared can be accessed by all threads in the block. Each block has a limited amount of shared (on-chip) memory.

Private can be accessed only by the thread itself. Each thread is allocated a limited number of registers.

The global memory is too slow compared to the shared and private memories. Therefore, the shared and private memory have to be utilized to the maximum extent.

2.3 Genetic Algorithms

One of the most suitable heuristics framework (*Meta Heuristics*) that exploits the availability of many cores performing the same instruction thread on multiple data is *Genetic Algorithms* (GA) [29]. Recently, efficient GPU-based GA are being proposed for solving hard problems. For example, Kang et. al. have introduced such a solution for the Traveller Salesman Problem (TSP) in [22]. We have not encountered any GPU-based GA for the PMP; even though, GA are recognized as one of the most effective evolutionary technique for solving optimization problems.

In GA, a large number (*Population*) of *Chromosomes* are generated and operated upon using similar operations like *mutations*, *crossover, migration* and *fitness test*. A chromosome is a finite sequence of genes commonly represented by a binary string or a set of integers. A crossover operation involves two parent chromosomes exchanging genes to produce offsprings, while a mutation involves only a single chromosome that mutates into a new one. Usually, each chromosome represents a candidate solution to the optimization problem, and the fitness of the chromosome represents the solution quality or the objective function value.

There exists a number of GA for solving the PMP in the literature, examples are [6][10][21]. Combining GA with local search method results in hybrid GA that could speed up the reach for global optima [13]. Hybrid GA for solving the PMP appear in [32][35]. Hybrid GA based on the variable neighborhood search have been recently published in [12][37]. Different hybrid GA for the PMP using solution archiving, greedy strategy, and fine-grained tournament selection are presented in [9][24] and [36], respectively.

3 A New Genetic Algorithm for the PMP Based on GPU and Pseudo Boolean Formulation

The algorithm is quite simple. The host randomly generates an initial population of chromosomes (candidate solutions) and passes them to a device kernel for fitness evaluations and enhancements. This basically iterates with migrating the best fit chromosomes from the current population to the next. The algorithm terminates at reaching an iteration limit or over saturating the solution enhancement.

Parallelism manifests in our algorithm in two folds. First, the host generates the next population in parallel with the device processing of the current population fitness evaluations and enhancements. Second, the kernel threads run in parallel, and each of which is assigned a chromosome for fitness evaluation and enhancement. The fitness evaluation is based on $\prod^{'}$ and $\triangle^{'}$ whose matrix structure harnessed the PMP data parallelism potentials as explained in 4.2. A chromosome enhancement is based on crossover and mutation operations. The details of these operations are explained in Subsections 4.3 and 4.4.

The following two subsections outline the algorithm host and device codes.

3.1 Host Code

n: Number of Clients,
m: Number of Facility Locations,
p: Number of Open Facilities,
C: $n \times m$ Distance/Cost Matrix
NB: Number of GPU Blocks
NT: Number of Threads per Block

EvolveLimit: Limit on Number of Calls to *Evolve* Kernel
S: Saturation Limit

Steps:

1. Read input *n, m, p, C, NB, NT, KernelsLimit*.
2. Call kernel (*Init* <<< *NB, NT* >>>) to differently seed *curand* in each thread.
3. Compute $\prod$ and $\triangle$ matrices each of size $(n \times m)$, and reduce them to $\prod^{'}$ and $\triangle^{'}$ each of size $(n \times m - p + 1)$.
4. Allocate device memory for $\prod^{'}$ and $\triangle^{'}$ using *cudaMalloc*.
5. Copy $\prod^{'}$ and $\triangle^{'}$ to device memory using *cudaMemcpy*.
6. Initialize a counter for the Number of *Evolve* kernal calls: NKernels = 0.
7. Generate candidate solutions as a random *population* of $NB \times NT$ *chromosomes*.
8. Wait for *Init* kernel to finish (*cudaDeviceSynchronize*).
9. Copy current population to device memory (*cudaMemcpy*).
10. Call kernel (*Evolve* <<< *NB, NT* >>>).
11. Generate a new random population of $NB \times NT$ *chromosomes* for next *Evolve* kernel call.
12. Wait for *Evolve* kernel to finish (*cudaDeviceSynchronize*).
13. Copy *most fit chromosome* (best solution) of each block to the Host memory and find *their most fit*.
14. Increment NKernels.
15. If (the *most fit chromosome* has not been changed by the last S *Evolve* kernal call) or (NKernels >= *EvolveLimit*), report the *most fit chromosome* as the best solution and stop;
16. Else *Migrate NB* best chromosomes to next population and *Go to* Step 9.
17. END (Host Code).

3.2 Device Code (Executed in Parallel by Each Thread)

3.2.1 Init Kernel

Steps:

1. Call *curand_ init(0, TID, 0, &state)*,
   (Thread Global Identifier) is passed as a *sequence number* [3]. This insures each thread to have a different random sequence when calling *curand*.
2. END (Init Kernel).

3.2.2 Evolve Kernel

Global Memory:

- Array B_MinCost[NB]: Minimum cost (Highest fitness) found by each block.
- Array Best[NB]: Chromosome with best fitness value for each block.

Shared Memory:

- Array MinCost[NT] : Minimum cost found by each thread in the block, initially set to MinCost[0].
- txMin[NT]: used to find the Index of the thread with best fitness value in the block, initially set to threadIdx.x.

Steps:

1. Evaluate the fitness of the thread chromosome $\mathbb{C}$ using $\prod^{'}$ and $\triangle^{'}$ matrices.
2. Initialize relative thread index and relative block size: rtx = threadIdx.x; rb_size = NT.
3. Crossover Cycle:
   For (CStride = NT /2 ; CStride >0; CStride = CStride/2)
   (a) Generate a random crossover point r1 (using *curand*).
   (b) If (rtx >= rb_size/2), Stride = -CStride else Stride = CStride.
   (c) Form the parent couples by finding a unique couple for each thread: Couple = TID + Stride;
   (d) Make *cross-over* between $\mathbb{C}$ and Couple at r1 and form offspring $\mathcal{F}$.
   (e) If (Fitness($\mathcal{F}$) < Fitness($\mathbb{C}$)), replace $\mathbb{C}$ by $\mathcal{F}$.
   (f) rb_size = rb_size /2; rtx = rtx % rb_size;
4. Synchronize each block threads using *syncthreads()*.

5. Mutation Cycle:
   For (i= lg(NT), Enhanced = False; (i > 0) and (Not Enhanced); i = i-1)
   (a) Randomly decide the mutation parameters as per the details explained in 4.4.
   (b) Mutate $\mathbb{C}$ to offspring $\mathcal{F}$.
   (c) If Fitness($\mathcal{F}$) < Fitness($\mathbb{C}$),
      i. Replace $\mathbb{C}$ by $\mathcal{F}$.
      ii. Enhanced = True.
6. Synchronize each block threads using *syncthreads()*.
7. Find best fitness in the block:
   For (Stride = NT /2 ; Stride >0; Stride = Stride/2)
   (a) tx = threadIdx.x.
   (b) If (tx < Stride and (MinCost[tx+Stride] < MinCost[tx]))
      i. MinCost[tx] =MinCost[tx+Stride].
      ii. txMin[tx]=txMin[tx+Stride].
8. If (threadIdx = 0), Store each block best fitness cost and its corresponding chromosome in B_MinCost[blockIdx.x] and Best[blockIdx.x], respectively.
9. End (Evolve Kernel).

4 Implementation Details

4.1 Chromosome Representation and Generation

A chromosome is represented as a vector $\mathbb{C}$ of m bits $\mathbb{C}_0$:$\mathbb{C}_{m-1}$, where *true* denotes an open facility and *false* denotes a closed one. Our algorithm generates chromosomes by random selection from a lexicographical order of a combinatorial sequence [17]. For each chromosome, it first generates a non-negative integer $i < \binom{m}{p}$ using a 64-bit random function. Next, it generates the i^{th} Lexicographic combination of $\binom{m}{p}$ using the efficient method presented in [28]. As a result, the generated chromosome will have exactly p true bits each of which corresponds to a selected element in the i^{th} Lexicographic combination of $\binom{m}{p}$. This method reduces the random function calls to one call per chromosome generation. Hence, it could maintain better quality random number generation as it limits the probability of exhaustively consuming the pseudo random sequence generated by the utilized random function. It has positively impacted the quality of the solutions generated by our algorithm. Interestingly, we are not aware of its existence in the literature.

4.2 Fitness Function

The Fitness function is a performance bottleneck. Each thread calls it several times. It is called to evaluate a thread assigned chromosome. It is also called in each crossover iteration to evaluate the offsprings. Furthermore, it is called to evaluate the mutation offsprings. Therefore, we designed this function to be as efficient as possible. In general, we harnessed the data parallelism in the PMP by using the pseudo boolean formulation and tailored it to be GPU suitable. We designed the fitness function to use $\prod'$ and $\triangle'$ rather than HBP. This enabled higher degree of data parallelism and restricted the required operations to be simple integer additions. We also took advantage of memory caching when accessing the $\prod'$ and $\triangle'$ simply because all threads in all blocks use exactly the same $\prod'$ and $\triangle'$ in read only mode.

The function algorithm is straightforward. For an input chromosome $\mathbb{C}$, it scans the entry of each client in $\triangle'$ as per the order of $\prod'$ and accumulates the corresponding increments till an open facility is found ($\mathbb{C}_{\mathbb{S}} = true$). The total accumulations of all clients represents the fitness of $\mathbb{C}$.

function Fitness ($\mathbb{C}$: Chromosome)

1. fitness = 0;
2. *for each* client i , $1 \leq i \leq$ n:
   - S =0;
   - *repeat*: S++; fitness = fitness + $\delta_{i\pi_{is}}$ *until* $\mathbb{C}_{\mathbb{S}}$
3. End (Fitness).

The Fitness function time order is $O(n(m-p))$ since each of $\prod'$ and $\triangle'$ is of size $n \times (m-p+1)$. In average, there will be one open facility in each $\frac{m-p-1}{p} \approx \frac{m}{p}$ locations assuming the p open facilities are normally distributed. Therefore, the Fitness function expected average time is $O(\frac{nm}{p})$. Based on this, we decided to compute each chromosome fitness by a single thread accumulating the increments of all clients. This results in a better device core utilization and higher scalability as explained in Section 6. Consequently, we designed the Evolve kernel grid and blocks to be of single dimensions.

4.3 Crossover Operation

Each thread determines its unique couple as in Step 3 of the Evolve kernel. One of the couple threads generates two random integers and shares them with its couple. The first integer $r1$, $0 \leq r1 < m$, determines the crossover starting index. The second integer $r2 = 2i$, $0 < i \leq \lfloor \frac{p}{2} \rfloor$, determines the number of genes to be exchanged. The exchanges count and occur only between unequal corresponding genes. In order to keep exactly p true genes in the offspring, exactly i genes are exchanged from 0 to 1, and the other remaining i genes are

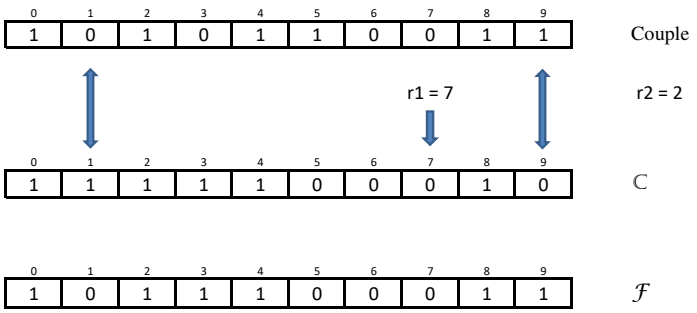

Figure 1: A crossover operation between $\mathbb{C}$ and Couple to offspring $\mathcal{F}$, $m = 10$, $p = 6$, $r1 = 7$, and $r2 = 2$.

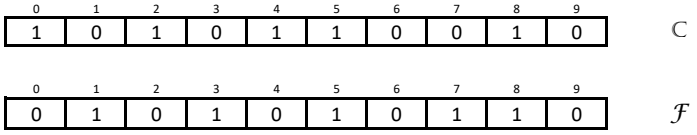

Figure 2: Three positions right circular shift mutation of $\mathbb{C}$ to offspring $\mathcal{F}$.

exchanged from 1 to 0. If the end of the chromosome is reached before having the right number of exchanges, the search continues from the beginning. If the the number of exchanges cannot reach $r2$, the operation fails and no offspring is produced. Figure 1 shows a crossover operation example.

4.4 Mutation Operation

The mutation operation is based on gene/bit shifting. We use two types of shifts: *circular shift* and *block shift*. In circular shift the number of genes to be shifted and the shift direction are randomly decided. Then, the genes are rotated accordingly. Figure 2 illustrates a three positions right circular shift mutation of chromosome $\mathbb{C}$ to offspring $\mathcal{F}$.

A block shift is a circular shift on a randomly selected subsequence of the chromosome to be mutated. The number of positions to be shifted, the shift direction, and the subsequence indexes are randomly decided. Then, the subsequence genes are rotated accordingly. Figure 3 shows one position left block shift on subsequence 3 to 6 of chromosome $\mathbb{C}$ to offspring $\mathcal{F}$.

4.5 Migration Operation

Cuda does not support thread synchronizations across different blocks. It only supports synchronizations of threads within the same block. Therefore, we had to implement the migration operation in the Host code. The migration takes place by selecting the best fit chromosome computed by each block in the last Evolve kernel and adding it to the newly generated population for the next

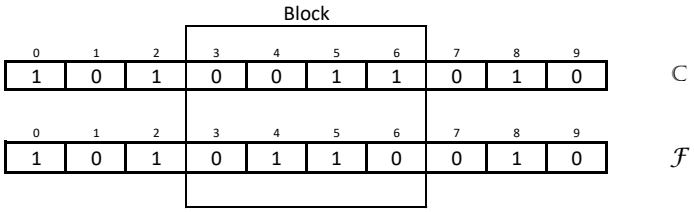

Figure 3: One position left block shift mutation of $\mathbb{C}$ to offspring $\mathcal{F}$.

kernel launch. Different variations of the migration operation can be applied. For example, the best of each block can be migrated to the same block in the next generation, or a team of all the bests can be migrated to a single block.

5 Time Complexity

The Fitness function time complexity as indicated in Subsection 4.2 is $O(n(m-p)) = O(nm)$. Hence, each of the crossover and mutation cycles in the Evolve kernel is $O(nm\ lg(NT))$. Therefore, the Evolve kernel time complexity is

$$T_E = O(nm\ lg(NT)).$$

The Init kernel is $O(1)$ since each thread would execute a constant number of operations.

The complexity of generating one chromosome (candidate solution) in the host is $O(mp)$. Consequently, the time complexity of the first eight steps in the Host Code is

$$T_S\ =\ O(nm) + O(NB\ NT\ mp).$$

The remaining steps are iterative and the time complexity of a single iteration of these steps is

$$T_H = O(NB\ NT\ mp)\ +\ W,$$

where W is the waiting time for the Evolve kernel to finish (Step 12). Since T_H and T_E run in parallel, the algorithm total complexity is

$$T_S\ + O(EvolveLimit\ \times\ max(T_H,\ T_E)).$$

For maximum utilization of both Host and GPU device, W must be 0 and $T_H\ =\ T_E$. This can be achieved by proper selection of NB and NT within the GPU device limits. Further synchronization between T_H and T_E could be achieved by increasing the crossover and/or mutation iterations to a number decided by an input parameter.

6 Experimentation Results

The objective of the experimentation was to test the effectiveness of our algorithm rather than to optimize the implementation to the best possible performance. We have tested the algorithm on all the benchmark instances we had access to. In total, we have tested 290 diversified instances collected as follows: 40 instances from the OR-Library [8], 40 instances of the so-called *complex instances* introduced in Table 2.6 of [16], and 210 instances from the discrete location problems benchmark library [2]. All our experimentation were executed on a Tesla K40 (2880 Cuda cores) hosted by HP Z820 workstation equipped with: 2× Intel Xeon processors 12 cores each, 16 GB RAM, and 2×512 GB solid state drives. The specifications of these equipment could be referred to in [19][31].

The algorithm succeeded in obtaining optimal solutions for all the 290 instances except two, namely OR-Library pmed30 and pmed40 where a better than 99.9% approximation was obtained for each. The obtained results are listed in Tables 1 to 9. By examining these results, we can draw the following notes and observations:

1. Our algorithm succeeded in obtaining optimal solutions for all what so called "complex instances" as shown in Table 9. Goldengorin et. al. introduced these forty instances where optimal solutions for thirty of which could not be obtained by linear programming using Elloumi formulation or pseudo boolean formulation and data reductions [16].

2. Our algorithm critically relies on randomization in initializing potential solutions and enhancing them. Thus, there is no guarantee to obtain the same results in each run of the algorithm on a given benchmark instance. Except for OR-Library pmed30 and pmed40, nevertheless, our implementation has shown excellent consistency in obtaining optimal solutions over multiple runs on each tested benchmark instance, but possibly with different timings and/or kernel counters. The measurements listed in Tables 1 to 9 are the medians obtained from different runs. Furthermore, these measurements are for the kernels in which the optimal solutions were achieved rather than for the kernels at which the program terminated with the exceptions of pmed30 and pmed40 as no optimal solutions were achieved.

3. The chromosomes generation method as explained in Subsection 4.1 tremendously enhances the candidate solutions' qualities. Moreover, the independence of the random functions used in the Host Code and in the Evolve kernel contributes to this enhancement as it dedicates the host random function for generating candidate solutions.

4. As indicated in Subsection 4.2, the Fitness function is a performance bottleneck. In average, there is one open facility in any $\frac{m-p-1}{p} \approx \frac{m}{p}$ locations assuming the open facilities are normally distributed. We could have sped up this function execution by using n threads each of which accumulates

the increments of one client (n threads scenario) rather than using a single thread to accumulate the increments of all clients (single thread scenario). In the n threads scenario, however, the Fitness function execution time is determined by the last finishing thread,(t_l), whose execution time in average will be worse than $\frac{m}{p}$ and could be $\Omega(m)$. Each of the other $n-1$ threads will be idle from its finishing time till the finishing time of t_l. This would result in underutilized device cores, and would hinder the performance scalability as n and m increase. The single thread scenario requires more time to evaluate the fitness of a single chromosome, but with no thread idle time. In this scenario, the Fitness function average time is $O(\frac{nm}{p})$. As p scales to $\theta(m)$, the average time could drop to $O(n)$. This explains the total time drop when scaling p and fixing n and m in our experimentation of the Pmed and the Complex benchmarks, refer to Tables 1 and 9. Evaluating the fitness of n chromosomes in the n threads scenario requires n^2 threads and $O(m)$ average time. The same requires n threads and $O(n)$ average time with proper scaling of p. As the number of threads exceeds the number of available cores, thread queuing overhead and waiting times will accumulate. Obviously, the n threads scenario requires more threads as n scales. Thus, it is more vulnerable to these overheads and waiting times.

5. The time needed to generate a population is proportional to its size = $NB \times NT$. We noticed that increasing the population size improves the chances of obtaining an optimal solution, refer to Table 6 as an example. However, determining the population size has to be within the GPU device hardware limitations: number of cores, threads queue depth, memory transfers, memory access conflicts, ... etc.

6. Increasing NT enhances the solutions' qualities as it increases the Crossover and Mutation iterations. This could result in obtaining an optimal solution in less number of Evolve kernel calls, but with more time per kernel. For example, compare Tables 2 and 3.

7. We have experimented the Crossover and Mutation impacts independently from NT. In these experimentation, we determined the number of Crossover and Mutations iterations by an input parameter. We found that increasing the number of iterations lead to optimal solutions in less number of kernels, but with higher average kernel time. Tables 4 and 7 show the related results.

8. The Migration operation impact starts from the second Evolve kernel and onward. As per our algorithm design, the number of chromosomes to be migrated to a next population is proportional to NB. Our experimentation indicated that the number of chromosomes to be migrated from each block and their distribution over the next kernel blocks influence the solution quality obtained by that kernel. We experienced these impacts while testing the Pmed, Chess Board, and Large Duality Gap-C benchmarks

as they required more kernel calls than the other benchmarks, refer to Tables 1, 3, and 8.

9. The experimentation results are consistent with the time complexity analysis in Section 5 except for the results shown in Tables 4 and 7 as explained above. The Evolve kernel average time is proportional to n, m, and TN. Furthermore, the experimentation indicated that this average time is also proportional to $\frac{TB \times TN}{Number\ of\ Cores}$. This is valid because threads will be queued as the number of threads exceeds the number of available cores. Tables 2 and 5 show the impact of increasing n and m on the kernel time when fixing NT and NB, while Tables 2 and 6 point out the impact of increasing $\frac{TB \times TN}{Number\ of\ Cores}$.

7 Conclusions

In this paper, we present a new genetic algorithm for the PMP based on GPU and pseudo-Boolean formulation. The algorithm is composed of Host code and Device code. The host randomly generates a population of chromosomes (candidate solutions) and passes them to a device kernel for fitness evaluations and enhancements. This basically iterates with migrating the best fit chromosomes from the current population to the next. The algorithm terminates at reaching an iteration limit or over saturating the solution enhancement. The algorithm is implemented using Cuda C version 7.5, and it was tested on 290 different benchmark instances. It has succeeded in obtaining optimal solutions for all the 290 instances except two for which better than 99.9% approximations have been obtained.

There are several venues for our future work on this topic. First, we will be working on identifying and developing solution enhancement operations that shall improve our algorithm performance. Second, we shall analyze and experiment the algorithm scalability limits on different GPGPU platforms. Third, we will investigate applying the presented algorithm on different variations of the facility location problems.

References

[1] *Cuda-C Programming Guide,* `http://docs.nvidia.com/cuda/cuda-c-programming-guide`.

[2] *Discrete location Problems Benchmark library, The P-median Problem, www.math.nsc.ru/AP/benchmarks/P-median.*

[3] *CURAND LIBRARY Programming Guide.* NVIDIA, September 2015.

[4] Bader AlBdaiwi, Diptesh Ghosh, and Boris Glodengorin. Data aggregation for p-median problems. *Journal of Combinatorial Optimization*, 21:348–363, 2011.

Instance Code	n = m	p	Number of Potential Solutions	Obtained Solution Approximation Ratio	Number of Kernel Calls	Time (Sec.)
Pmed 1	100	5	7.53E+07	Optimal	1	2
Pmed 2	100	10	1.73E+13	Optimal	1	2
Pmed 3	100	10	1.73E+13	Optimal	1	2
Pmed 4	100	20	5.36E+20	Optimal	1	2
Pmed 5	100	33	2.95E+26	Optimal	2	4
Pmed 6	200	5	2.54E+09	Optimal	1	8
Pmed 7	200	10	2.25E+16	Optimal	1	6
Pmed 8	200	20	1.61E+27	Optimal	4	15
Pmed 9	200	40	2.05E+42	Optimal	7	23
Pmed 10	200	67	1.45E+54	Optimal	13	49
Pmed 11	300	5	1.96E+10	Optimal	1	13
Pmed 12	300	10	1.40E+18	Optimal	2	19
Pmed 13	300	30	1.73E+41	Optimal	10	61
Pmed 14	300	60	9.04E+63	Optimal	14	137
Pmed 15	300	100	4.16E+81	Optimal	16	3744
Pmed 16	400	5	8.32E+10	Optimal	1	23
Pmed 17	400	10	2.58E+19	Optimal	4	61
Pmed 18	400	40	1.97E+55	Optimal	15	158
Pmed 19	400	80	4.23E+85	Optimal	15	2608
Pmed 20	400	133	1.26E+109	Optimal	17	462
Pmed 21	500	5	2.55E+11	Optimal	1	34
Pmed 22	500	10	2.46E+20	Optimal	6	150
Pmed 23	500	50	2.31E+69	Optimal	27	495
Pmed 24	500	100	2.04E+107	Optimal	16	4104
Pmed 25	500	167	7.85E+136	Optimal	14	2201
Pmed 26	600	5	6.37E+11	Optimal	1	50
Pmed 27	600	10	1.55E+21	Optimal	9	280
Pmed 28	600	60	2.77E+83	Optimal	13	2918
Pmed 29	600	120	1.01E+129	Optimal	50	8856
Pmed 30	600	200	2.51E+164	0.999497487	100	41687
Pmed 31	700	5	1.38E+12	Optimal	1	98
Pmed 32	700	10	7.30E+21	Optimal	3	186
Pmed 33	700	70	3.37E+97	Optimal	28	5385
Pmed 34	700	140	5.03E+150	Optimal	33	8143
Pmed 35	800	5	2.70E+12	Optimal	2	249
Pmed 36	800	10	2.80E+22	Optimal	2	163
Pmed 37	800	80	4.14E+111	Optimal	15	8326
Pmed 38	900	5	4.87E+12	Optimal	5	763
Pmed 39	900	10	9.14E+22	Optimal	7	643
Pmed 40	900	90	5.13E+125	0.99980503	100	63088

Table 1: Results for P median benchmark instances obtained from OR-Library with NB = 60 and NT = 256.

Instance Code	n = m	p	Number of Potential Solutions	Obtained Solution Approximation Ratio	Number of Kernel Calls	Time (Sec.)
313	128	16	9.3343E+19	Optimal	1	7
323	128	16	9.3343E+19	Optimal	1	7
334	128	16	9.3343E+19	Optimal	1	7
434	128	16	9.3343E+19	Optimal	1	7
534	128	16	9.3343E+19	Optimal	1	7
634	128	16	9.3343E+19	Optimal	2	12
734	128	16	9.3343E+19	Optimal	1	7
834	128	16	9.3343E+19	Optimal	1	7
934	128	16	9.3343E+19	Optimal	1	7
1034	128	16	9.3343E+19	Optimal	1	7
1134	128	16	9.3343E+19	Optimal	1	7
1234	128	16	9.3343E+19	Optimal	1	6
1334	128	16	9.3343E+19	Optimal	1	7
1434	128	16	9.3343E+19	Optimal	1	7
1534	128	16	9.3343E+19	Optimal	1	7
1634	128	16	9.3343E+19	Optimal	1	6
1734	128	16	9.3343E+19	Optimal	2	13
1834	128	16	9.3343E+19	Optimal	1	7
1934	128	16	9.3343E+19	Optimal	1	7
2034	128	16	9.3343E+19	Optimal	1	7
2134	128	16	9.3343E+19	Optimal	1	7
2234	128	16	9.3343E+19	Optimal	1	7
2334	128	16	9.3343E+19	Optimal	1	7
2434	128	16	9.3343E+19	Optimal	1	7
2534	128	16	9.3343E+19	Optimal	1	7
2634	128	16	9.3343E+19	Optimal	1	7
2734	128	16	9.3343E+19	Optimal	1	7
2834	128	16	9.3343E+19	Optimal	1	7
2934	128	16	9.3343E+19	Optimal	1	7
3034	128	16	9.3343E+19	Optimal	1	7

Table 2: Results for Perfect Codes Instances obtained from Discrete Location Problems Benchmark Library with NB = 120 and NT = 256.

Instance Code	n = m	p	Number of Potential Solutions	Obtained Solution Approximation Ratio	Number of Kernel Calls	Time (Sec.)
334	144	16	6.88E+20	Optimal	11	36
434	144	16	6.88E+20	Optimal	12	39
534	144	16	6.88E+20	Optimal	11	59
634	144	16	6.88E+20	Optimal	14	45
734	144	16	6.88E+20	Optimal	7	25
834	144	16	6.88E+20	Optimal	9	49
934	144	16	6.88E+20	Optimal	13	42
1034	144	16	6.88E+20	Optimal	29	100
1134	144	16	6.88E+20	Optimal	7	24
1234	144	16	6.88E+20	Optimal	21	67
1334	144	16	6.88E+20	Optimal	8	29
1434	144	16	6.88E+20	Optimal	10	34
1534	144	16	6.88E+20	Optimal	19	5
1634	144	16	6.88E+20	Optimal	18	5
1734	144	16	6.88E+20	Optimal	18	56
1834	144	16	6.88E+20	Optimal	46	142
1934	144	16	6.88E+20	Optimal	8	44
2034	144	16	6.88E+20	Optimal	11	37
2134	144	16	6.88E+20	Optimal	14	44
2234	144	16	6.88E+20	Optimal	22	75
2334	144	16	6.88E+20	Optimal	13	69
2434	144	16	6.88E+20	Optimal	7	23
2534	144	16	6.88E+20	Optimal	11	35
2634	144	16	6.88E+20	Optimal	2	9
2734	144	16	6.88E+20	Optimal	6	23
2834	144	16	6.88E+20	Optimal	6	20
2934	144	16	6.88E+20	Optimal	1	6
3034	144	16	6.88E+20	Optimal	17	54
3134	144	16	6.88E+20	Optimal	11	37
3234	144	16	6.88E+20	Optimal	10	33

Table 3: Results for Chess Board Instances obtained from Discrete Location Problems Benchmark Library with NB = 480 and NT = 96.

Instance Code	n = m	p	Number of Potential Solutions	Obtained Solution Approximation Ratio	Number of Kernel Calls	Time (Sec.)
1	133	12	3.84E+16	Optimal	1	78
2	133	12	3.84E+16	Optimal	1	78
3	133	12	3.84E+16	Optimal	1	81
4	133	12	3.84E+16	Optimal	1	78
5	133	12	3.84E+16	Optimal	1	76
6	133	12	3.84E+16	Optimal	1	79
7	133	12	3.84E+16	Optimal	1	75
8	133	12	3.84E+16	Optimal	1	76
9	133	12	3.84E+16	Optimal	1	76
10	133	12	3.84E+16	Optimal	1	78
11	133	12	3.84E+16	Optimal	1	75
12	133	12	3.84E+16	Optimal	1	76
13	133	12	3.84E+16	Optimal	1	76
14	133	12	3.84E+16	Optimal	1	81
15	133	12	3.84E+16	Optimal	1	76
16	133	12	3.84E+16	Optimal	1	77
17	133	12	3.84E+16	Optimal	1	75
18	133	12	3.84E+16	Optimal	1	80
19	133	12	3.84E+16	Optimal	1	75
20	133	12	3.84E+16	Optimal	1	77
21	133	12	3.84E+16	Optimal	1	77
22	133	12	3.84E+16	Optimal	1	75
23	133	12	3.84E+16	Optimal	1	77
24	133	12	3.84E+16	Optimal	1	74
25	133	12	3.84E+16	Optimal	1	78
26	133	12	3.84E+16	Optimal	1	74
27	133	12	3.84E+16	Optimal	1	78
28	133	12	3.84E+16	Optimal	1	79
29	133	12	3.84E+16	Optimal	1	73
30	133	12	3.84E+16	Optimal	1	76

Table 4: Results for Finite Projective Planes Instances, K = 11 obtained from Discrete Location Problems Benchmark Library with NB = 120 and NT = 256. In this experiment, the number of crossover and mutation iterations were preset to 60.

Instance Code	n = m	p	Number of Potential Solutions	Obtained Solution Approximation Ratio	Number of Kernel Calls	Time (Sec.)
1	307	18	5.51E+28	Optimal	1	39
2	307	18	5.51E+28	Optimal	1	40
3	307	18	5.51E+28	Optimal	1	40
4	307	18	5.51E+28	Optimal	1	40
5	307	18	5.51E+28	Optimal	1	40
6	307	18	5.51E+28	Optimal	6	212
7	307	18	5.51E+28	Optimal	1	39
8	307	18	5.51E+28	Optimal	1	39
9	307	18	5.51E+28	Optimal	1	39
10	307	18	5.51E+28	Optimal	1	39
11	307	18	5.51E+28	Optimal	1	38
12	307	18	5.51E+28	Optimal	2	73
13	307	18	5.51E+28	Optimal	2	73
14	307	18	5.51E+28	Optimal	1	39
15	307	18	5.51E+28	Optimal	2	74
16	307	18	5.51E+28	Optimal	2	71
17	307	18	5.51E+28	Optimal	1	39
18	307	18	5.51E+28	Optimal	1	39
19	307	18	5.51E+28	Optimal	1	39
20	307	18	5.51E+28	Optimal	1	40
21	307	18	5.51E+28	Optimal	1	39
22	307	18	5.51E+28	Optimal	5	178
23	307	18	5.51E+28	Optimal	4	140
24	307	18	5.51E+28	Optimal	1	40
25	307	18	5.51E+28	Optimal	1	39
26	307	18	5.51E+28	Optimal	1	38
27	307	18	5.51E+28	Optimal	2	75
28	307	18	5.51E+28	Optimal	1	40
29	307	18	5.51E+28	Optimal	1	41
30	307	18	5.51E+28	Optimal	1	39

Table 5: Results for Finite Projective Planes Instances, K = 17 obtained from Discrete Location Problems Benchmark Library with NB = 120 and NT = 256.

Instance Code	n = m	p	Number of Potential Solutions	Obtained Solution Approximation Ratio	Number of Kernel Calls	Time (Sec.)
332	100	12	1.05E+15	Optimal	1	46
432	100	12	1.05E+15	Optimal	2	85
532	100	12	1.05E+15	Optimal	2	86
632	100	12	1.05E+15	Optimal	1	46
732	100	12	1.05E+15	Optimal	1	46
832	100	12	1.05E+15	Optimal	1	45
932	100	12	1.05E+15	Optimal	1	45
1032	100	12	1.05E+15	Optimal	1	60
1132	100	12	1.05E+15	Optimal	1	60
1232	100	12	1.05E+15	Optimal	1	52
1332	100	12	1.05E+15	Optimal	1	46
1432	100	12	1.05E+15	Optimal	1	47
1532	100	12	1.05E+15	Optimal	2	91
1632	100	12	1.05E+15	Optimal	1	48
1732	100	12	1.05E+15	Optimal	1	49
1832	100	12	1.05E+15	Optimal	7	304
1932	100	12	1.05E+15	Optimal	1	48
2032	100	12	1.05E+15	Optimal	1	48
2132	100	12	1.05E+15	Optimal	1	49
2232	100	12	1.05E+15	Optimal	1	49
2332	100	12	1.05E+15	Optimal	1	49
2432	100	12	1.05E+15	Optimal	1	49
2532	100	12	1.05E+15	Optimal	1	49
2632	100	12	1.05E+15	Optimal	2	91
2732	100	12	1.05E+15	Optimal	25	1027
2832	100	12	1.05E+15	Optimal	1	49
2932	100	12	1.05E+15	Optimal	1	45
3032	100	12	1.05E+15	Optimal	1	45
3132	100	12	1.05E+15	Optimal	1	43
3232	100	12	1.05E+15	Optimal	2	88

Table 6: Results for Large Duality Gap-A Instances obtained from Discrete Location Problems Benchmark Library with NB = 1500 and NT = 64.

Instance Code	n = m	p	Number of Potential Solutions	Obtained Solution Approximation Ratio	Number of Kernel Calls	Time (Sec.)
331	100	14	4.42E+16	Optimal	2	54
431	100	14	4.42E+16	Optimal	2	54
531	100	14	4.42E+16	Optimal	2	52
631	100	14	4.42E+16	Optimal	2	58
731	100	14	4.42E+16	Optimal	5	130
831	100	14	4.42E+16	Optimal	1	27
931	100	14	4.42E+16	Optimal	2	58
1031	100	14	4.42E+16	Optimal	1	30
1131	100	14	4.42E+16	Optimal	1	30
1231	100	14	4.42E+16	Optimal	6	156
1331	100	14	4.42E+16	Optimal	1	29
1431	100	14	4.42E+16	Optimal	5	138
1531	100	14	4.42E+16	Optimal	6	171
1631	100	14	4.42E+16	Optimal	2	52
1731	100	14	4.42E+16	Optimal	1	30
1831	100	14	4.42E+16	Optimal	2	52
1931	100	14	4.42E+16	Optimal	1	26
2031	100	14	4.42E+16	Optimal	1	27
2131	100	14	4.42E+16	Optimal	1	26
2231	100	14	4.42E+16	Optimal	2	62
2331	100	14	4.42E+16	Optimal	1	28
2431	100	14	4.42E+16	Optimal	5	132
2531	100	14	4.42E+16	Optimal	1	27
2631	100	14	4.42E+16	Optimal	10	252
2731	100	14	4.42E+16	Optimal	2	61
2831	100	14	4.42E+16	Optimal	1	32
2931	100	14	4.42E+16	Optimal	15	387
3031	100	14	4.42E+16	Optimal	1	26
3131	100	14	4.42E+16	Optimal	1	30
3231	100	14	4.42E+16	Optimal	3	75

Table 7: Results for Large Duality Gap-B Instances obtained from Discrete Location Problems Benchmark Library with NB = 120 and NT = 256. In this experiment, the number of crossover and mutation iterations was preset to 20.

Instance Code	n = m	p	Number of Potential Solutions	Obtained Solution Approximation Ratio	Number of Kernel Calls	Time (Sec.)
333	100	14	4.42E+16	Optimal	3	37
433	100	14	4.42E+16	Optimal	2	26
533	100	14	4.42E+16	Optimal	2	26
633	100	14	4.42E+16	Optimal	4	47
733	100	14	4.42E+16	Optimal	3	38
833	100	14	4.42E+16	Optimal	3	37
933	100	14	4.42E+16	Optimal	2	26
1033	100	14	4.42E+16	Optimal	4	49
1133	100	14	4.42E+16	Optimal	7	80
1233	100	14	4.42E+16	Optimal	11	126
1333	100	14	4.42E+16	Optimal	20	218
1433	100	14	4.42E+16	Optimal	2	27
1533	100	14	4.42E+16	Optimal	32	363
1633	100	14	4.42E+16	Optimal	3	37
1733	100	14	4.42E+16	Optimal	2	26
1833	100	14	4.42E+16	Optimal	18	198
1933	100	14	4.42E+16	Optimal	3	36
2033	100	14	4.42E+16	Optimal	5	59
2133	100	14	4.42E+16	Optimal	7	81
2233	100	14	4.42E+16	Optimal	2	26
2333	100	14	4.42E+16	Optimal	4	47
2433	100	14	4.42E+16	Optimal	9	105
2533	100	14	4.42E+16	Optimal	22	245
2633	100	14	4.42E+16	Optimal	3	37
2733	100	14	4.42E+16	Optimal	78	873
2833	100	14	4.42E+16	Optimal	2	27
2933	100	14	4.42E+16	Optimal	8	92
3033	100	14	4.42E+16	Optimal	1	15
3133	100	14	4.42E+16	Optimal	2	26
3233	100	14	4.42E+16	Optimal	3	37

Table 8: Results for Large Duality Gap-C Instances obtained from Discrete Location Problems Benchmark Library with NB = 1500 and NT = 32.

n = m	p	Number of Potential Solutions	Obtained Solution Approximation Ratio	Number of Kernel Calls	Time (Sec.)
100	5	7.53E+07	Optimal	1	10
100	10	1.73E+13	Optimal	1	11
100	20	5.36E+20	Optimal	1	11
100	33	2.95E+26	Optimal	1	9
200	5	2.54E+09	Optimal	1	41
200	10	2.25E+16	Optimal	1	39
200	20	1.61E+27	Optimal	1	40
200	40	2.05E+42	Optimal	1	62
200	67	1.45E+54	Optimal	1	50
300	5	1.96E+10	Optimal	1	86
300	10	1.40E+18	Optimal	1	143
300	30	1.73E+41	Optimal	1	83
300	60	9.04E+63	Optimal	1	95
300	100	4.16E+81	Optimal	1	111
400	5	8.32E+10	Optimal	1	256
400	10	2.58E+19	Optimal	1	253
400	40	1.97E+55	Optimal	1	147
400	80	4.23E+85	Optimal	1	172
400	133	1.26E+109	Optimal	1	197
500	5	2.55E+11	Optimal	1	401
500	10	2.46E+20	Optimal	1	389
500	50	2.31E+69	Optimal	2	430
500	100	2.04E+107	Optimal	2	486
500	167	7.85E+136	Optimal	1	313
600	5	6.37E+11	Optimal	1	579
600	10	1.55E+21	Optimal	2	604
600	60	2.77E+83	Optimal	2	617
600	120	1.01E+129	Optimal	2	706
600	200	2.51E+164	Optimal	1	453
700	5	1.38E+12	Optimal	2	846
700	10	7.30E+21	Optimal	2	818
700	70	3.37E+97	Optimal	2	839
700	140	5.03E+150	Optimal	2	960
800	5	2.70E+12	Optimal	3	1660
800	10	2.80E+22	Optimal	2	1507
800	80	4.14E+111	Optimal	2	1123
900	5	4.87E+12	Optimal	1	1291
900	10	9.14E+22	Optimal	2	1905
900	90	5.13E+125	Optimal	2	1407

Table 9: Results for the Complex Instances Introduced in [16], where NB = 120 and NT = 256.

[5] Bader F AlBdaiwi, Boris Goldengorin, and Gerard Sierksma. Equivalent instances of the simple plant location problem. *Computers & Mathematics with Applications*, 57(5):812–820, 2009.

[6] Osman Alp and Erhan Erkut. An efficient genetic algorithm for the p-median problem. *Annals of Oerations Research*, 122:21–42, 2003.

[7] Fabiano Fernandes Bargos, Wendell de Queiroz Lamas, Danubia Caporusso Bargos, Morun Bernardino Neto, and Paula Cristiane Pinto Mesquita Pardal. Location problem method applied to sugar and ethanol mills location optimization. *Renewable and Sustainable Energy Reviews*, 65:274–282, 2016.

[8] J. E. Beasley. *OR-LIBRARY, http://people.brunel.ac.uk/~mastjjb/jeb/orlib/pmedinfo.html*

[9] Benjamin Biesinger, Bin Hu, and Günther Raidl. A hybrid genetic algorithm with solution archive for the discrete (r| p)-centroid problem. *Journal of Heuristics*, 21(3):391–431, 2015.

[10] Burcin Bozkaya, Jianjun Zhang, and Erhan Erkut. An efficient genetic algorithm for the p-median problem. *Facility location: Applications and theory*, pages 179–205, 2002.

[11] Mark S. Daskin and Kayse Lee Maass. *Location Science, G. Laporte, S. Nickel and F. Saldanha da Gama (Eds.)*, chapter 2: The p-Median Problem, pages 21–45. Springer, 2015.

[12] Zvi Drezner, Jack Brimberg, Nenad Mladenović, and Said Salhi. New heuristic algorithms for solving the planar p-median problem. *Computers & Operations Research*, 62:296–304, 2015.

[13] Tarek A El-Mihoub, Adrian A Hopgood, Lars Nolle, and Alan Battersby. Hybrid genetic algorithms: A review. *Engineering Letters*, 13(2):124–137, 2006.

[14] Reza Zanjirani Farahani, Masoud Hekmatfar, Alireza Boloori Arabani, and Ehsan Nikbakhsh. Hub location problems: A review of models, classification, solution techniques, and applications. *Computers & Industrial Engineering*, 64(4):1096–1109, 2013.

[15] Boris Goldengorin, Anton Kocheturov, and Panos M Pardalos. A pseudo-boolean approach to the market graph analysis by means of the p-median model. In *Clusters, Orders, and Trees: Methods and Applications*, pages 77–89. Springer, 2014.

[16] Boris Goldengorin, Dmitry Krushinsky, and Panos Pardalos. *Cell Formation in Industrial Engineering, Theory, Algorithms and Experiments*. Springer, 2013.

[17] Marshall Hall and Donald E Knuth. Combinatorial analysis and computers. *The American Mathematical Monthly*, 72(2):21–28, 1965.

[18] Peter Ladislaw Hammer. Plant location - a pseudo-boolean approach. *Israel Journal of Technology*, 6(5):330–332, 1968.

[19] HP Inc., USA. *HP Z820 Workstation*, April 2015. c04111526 - DA - 14264 - Worldwide - Version 48.

[20] Patrick Jaillet, Gao Song, and Gang Yu. Airline network design and hub location problems. *Location science*, 4(3):195–212, 1996.

[21] Jorge H Jaramillo, Joy Bhadury, and Rajan Batta. On the use of genetic algorithms to solve location problems. *Computers & Operations Research*, 29(6):761–779, 2002.

[22] Semin Kang, Sung-Soo Kim, Jongho Won, and Young-Min Kang. Gpu-based parallel genetic approach to large-scale travelling salesman problem. *Journal of Super Computing*, pages 1–16, May 2016. DOI: 10.1007/s11227-016-1748-1.

[23] Oded Kariv and S Louis Hakimi. An algorithmic approach to network location problems. II: The p-medians. *SIAM Journal on Applied Mathematics*, 37(3):539–560, 1979.

[24] Lev Aleksandrovich Kazakovtsev, Victor Orlov, Aljona Aleksandrovna Stupina, and Vladimir Kazakovtsev. Modied genetic algorithm with greedy heuristic for continuous and discrete p-median problems. *Facta Universitatis, Series: Mathematics and Informatics*, 30(1):89–106, 2015.

[25] Gilbert Laporte, Stefan Nickel, and Francisco Saldanha da Gama. *Location science*. Springer, 2015.

[26] Gino Lim and Likang Ma. Gpu-based parallel vertex substitution algorithm for the p-median problem. *Computers & Industrial Engineering*, 64(1):381–388, January 2013.

[27] Likang Ma and Gino Lim. Gpu-based parallel computational algorithms for solving p-median problem. In *IIE Annual Conference. Proceedings*, page 1. Institute of Industrial Engineers-Publisher, 2011.

[28] James McCaffrey. Generating the mth lexicographical element of a mathematical combination. *CTAN: `http://msdn.microsoft.com/en-us/library/aa289166`*, July 2004.

[29] Melanie Mitchell. *An Introduction to Genetic Algorithms*. MIT Press, Cambridge, 1998.

[30] Nenad Mladenović, Jack Brimberg, Pierre Hansen, and José A Moreno-Pérez. The p-median problem: A survey of metaheuristic approaches. *European Journal of Operational Research*, 179(3):927–939, 2007.

[31] Nvidia Corporation, USA. *Tesla K40 Active Accelerator*, November 2013. BD-06949-001_v03.

[32] Pascal Rebreyend, Laurent Lemarchand, and Reinhardt Euler. A computational comparison of different algorithms for very large p-median problems. In *15^{th} European Conference on Evolutionary Computation in Combinatorial Optimization*, pages 13–24. Springer, 2015.

[33] Josh Reese. Solution methods for the p-median problem: An annotated bibliography. *Networks*, 48(3):125–142, 2006.

[34] Yonglin Ren and Anjali Awasthi. Investigating metaheuristics applications for capacitated location allocation problem on logistics networks. In *Chaos Modeling and Control Systems Design*, pages 213–238. Springer, 2015.

[35] Mauricio GC Resende and Renato F Werneck. A hybrid heuristic for the p-median problem. *Journal of heuristics*, 10(1):59–88, 2004.

[36] Zorica Stanimirović. A genetic algorithm approach for the capacitated single allocation p-hub median problem. *Computing and Informatics*, 29(1):117–132, 2012.

[37] Raca Todosijević, Dragan Urošević, Nenad Mladenović, and Saïd Hanafi. A general variable neighborhood search for solving the uncapacitated r-allocation p-hub median problem. *Optimization Letters*, pages 1–13, 2015.